

REMARKS

I. Introduction

In response to the Office Action dated July 26, 2007, which was made final, and in conjunction with the Request for Continued Examination (RCE) submitted herewith, claims 1, 4 and 7 have been amended. Claims 1-9 remain in the application. Re-examination and re-consideration of the application, as amended, is requested.

II. Prior Art Rejections

A. The Office Action Rejections

In paragraphs (2)-(4) of the Office Action, claims 1-9 were rejected under 35 U.S.C. §103(a) as being obvious in view of the combination of U.S. Patent 5,963,936 (Cochrane) in view of U.S. Patent No. 7,010,524 (Galindo-Legaria).

Applicants' attorney respectfully traverses these rejections in view of the amendments above and the arguments below.

B. The Applicants' Independent Claims

Independent claims 1, 4 and 7 are directed to a method, system and article of manufacture for optimizing a query. Claim 1 is representative and recites a method of optimizing a query in a computer system, the query being performed by the computer system to retrieve data from a database stored on the computer system, the method comprising: (a) during compilation of the query, maintaining a GROUP BY clause with one or more GROUPING SETS, ROLLUP or CUBE operations in its original form, instead of rewriting the GROUP BY clause, until after query rewrite; (b) at a later stage of query compilation, translating the GROUP BY clause with the GROUPING SETS, ROLLUP or CUBE operations into a plurality of levels, wherein each of the levels has one or more grouping sets comprised of grouping columns, and generating a query execution plan for the query with a super group block having an array of pointers, wherein each pointer points to the grouping sets for a particular one of the levels; and (c) performing the query execution plan to retrieve data from a database stored on the computer system.

C. The Cochrane Reference

Cochrane describes a method and apparatus for detecting and stacking grouping sets to support GROUP BY operations with GROUPING SETS, ROLLUP and CUBE extensions in

relational database management systems, with greatly reduced numbers of grouping sets. A first GROUP BY (element-list1) is input to a second GROUP BY (element-list2), resulting in the GROUP BY of the intersection of the two lists. This intersection property is then useable to reduce the number of GROUP BYs required to implement the grouping by GROUPING SETS, ROLLUPS, and CUBEs required for the online analytical processing of data contained in the database.

D. The Galindo-Legaria Reference

Galindo-Legaria describes validation of large numbers of alternative execution plans for a database query, either an exhaustive enumeration of the complete space of alternatives, or else an unbiased random sample, that is performed by efficiently constructing execution trees from a data structure having groups alternative operators that are ranked in a directory. Each global rank of a plan identifies that plan uniquely among all the alternative plans. The operators are unranked from the directory according to a specification that characterizes the desired plans.

E. Applicants' Claims Are Patentable Over The References

Applicants' invention, as recited in amended independent claims 1, 4 and 7, is patentable over the combination of Cochrane and Galindo-Legaria, because the claims recite limitations not found in the references. Specifically, Applicants' independent claims 1, 4 and 7 have been amended to better distinguish over the references.

Nonetheless, according to the Office Action, Cochrane reaches the "maintaining" element of Applicants' independent claims at col. 7, lines 26-30 and 44-48, and Cochrane teaches the "translating" element of Applicants' independent claims at col. 8, lines 26-42, Figure 7 and col. 11, line 62 – col. 12, line 15. However, the Office Action admits that Cochrane does not reach generating a query execution plan with a super group block having an array of grouping sets, wherein each pointer points to a linked list representing grouping sets for a particular level. Nonetheless, the Office Action asserts that Galindo-Legaria teaches these limitations at col. 5, lines 25-34, col. 5, lines 56-63 and Figure 3.

Moreover, in response to Applicants' arguments, the Office Action states the following:

Response to Arguments

4. Applicant's arguments filed 8 May 2006 have been fully considered but they are not persuasive.

Applicant argues that the combination of Cochrane et al. and Galindo-Legaria et al. does not maintain the GROUP BYs in their original form until after query rewrite.

Applicant argues that "instead, the optimization scheme of Cochrane et al. reduces the GROUP BYs during query rewrite, which necessarily comprises not maintaining the GROUP BY clause until after query rewrite". It is noted that once the query is reduced, it is 'rewritten'. As such, it was maintained in its original form until it was rewritten.

Applicant argues that replacing the groups of operators in Galindo-Legaria et al. with the grouping sets comprised of grouping columns in Cochrane et al. would render Galindo-Legaria et al. inoperative. However, Examiner notes that the GROUP BY command is an operator, and that Cochrane et al. teaches having a list of grouping sets in the form of GROUP BYs that are output to the query optimizer to develop a plan (see Figure 7 and 8:26-42).

Applicant also argues that the modification suggested by the Office Action is impermissible hindsight. In response to applicant's argument that the examiner's conclusion of obviousness is based upon improper hindsight reasoning, it must be recognized that any judgment on obviousness is in a sense necessarily a reconstruction based upon hindsight reasoning. But so long as it takes into account only knowledge which was within the level of ordinary skill at the time the claimed invention was made, and does not include knowledge gleaned only from the applicant's disclosure, such a reconstruction is proper. See *In re McLaughlin*, 443 F.2d 1392, 170 USPQ 209 (CCPA 1971). In this case, as Galindo-Legaria et al. teaches utilizing operators to generate a query execution plan, and that the GROUPING SETS and GROUP BYs of Cochrane et al. are operators, in addition to the motivation provided in the rejection above, it would have been obvious to one of ordinary skill in the art at the time the invention was made to have modified Cochrane et al. to include the teaching of Galindo-Legaria et al.

Applicants' attorney disagrees with this analysis.

The combination of Cochrane and Galindo-Legaria does not teach or suggest all the limitations recited in amended independent claims 1, 4 and 7.

With regard to the limitations "during compilation of the query, maintaining a GROUP BY clause with one or more GROUPING SETS, ROLLUP or CUBE operations in its original form, instead of rewriting the GROUP BY clause, until after query rewrite," and "at a later stage of query compilation, translating the GROUP BY clause with the GROUPING SETS, ROLLUP or CUBE operations into a plurality of levels, wherein each of the levels has one or more grouping sets comprised of grouping columns," these limitations are not shown in Cochrane at col. 7, lines 26-30, col. 7, lines 44-48, col. 8, lines 26-42, Figure 7, and col. 11, line 62 through col. 12, line 15, which are set forth below:

Cochrane: Col. 7, Lines 26-30

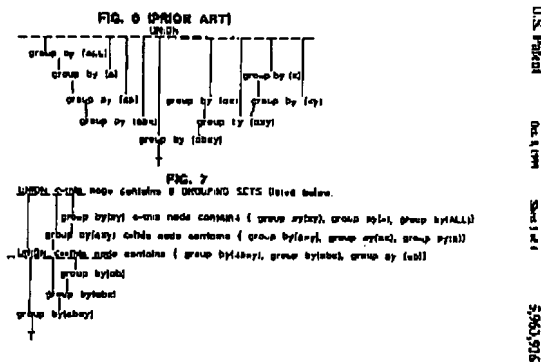
Generally, the query parser 92 lexes, parses, and semantically checks a query, producing an internal representation (a "query graph model") that is rewritten and submitted to the optimizer which generates an optimized query execution plan.

Cochrane: Col. 7, Lines 44-48

The system of FIG. 5 employs the invention to produce a QGM in which the number of GROUP BYs necessary to execute a GROUP BY with multiple GROUPING SETS, concatenated ROLLUPS, or a CUBE has been reduced.

Cochrane: Col. 8, Lines 26-42

Now, utilizing the principles of the present invention, and noting the previously derived intersection results shown above at (1)-(4), it becomes possible to construct a query graph model that includes a stacking of GROUP BYs that results in the computation and planning of only 5 GROUP BYs as opposed to the 9 required in FIG. 6. This query graph model is shown in FIG. 7. It should be emphasized that the query graph model of FIG. 7 produces results that are identical to the solution provided in FIG. 6, with only 5 GROUP BY operations, a considerable economy in computational overhead. Indeed, this reduction in the number of GROUP BYs may, in an RDBMS implementing large multi-dimensional tables and subject to complex OLAP queries, be necessary to implement the query. This is due to the fact that the size of such queries, combined with the prior art, can require such large-scale computational assets as to render the query incapable of implementation.

Cochrane: Figure 7Cochrane: Col. 11, line 62 – col. 12, line 15

As an example, consider the following: GROUP BY ROLLUP(a,b), ROLLUP(x,y) in which the GROUP BY's for ROLLUP(a,b) are:

GROUP BY(a,x,y)

GROUP BY(x,y)

and the GROUP BY's for ROLLUP(x,y) are:

GROUP BY(a,b,x)

GROUP BY(a,b)

Now, the base group for ROLLUP(a,b)ROLLUP(x,y) is determined by base step:

```

UNION1
|||
|| GROUP BY(x,y)
|||
|| GROUP BY(a,x,y)
|||
|| GROUP BY(a,b,x,y)
|
T

```

The description from Cochrane set forth above merely describes the translation of a query into a "query graph model" that is rewritten and submitted to an optimizer which generates an optimized query execution plan, wherein the optimization of GROUP BYs is performed by stacking, which reduces the number of GROUP BYs while producing identical results.

However, this optimization scheme of Cochrane says nothing about maintaining a GROUP BY clause with one or more GROUPING SETS, ROLLUP or CUBE operations in its original form, instead of rewriting the GROUP BY clause, until after query rewrite.

Instead, the optimization scheme of Cochrane reduces the GROUP BYs during query rewrite, which means that the GROUP BY clause is not maintained in its original form until after query rewrite, but instead the GROUP BY clause is rewritten.

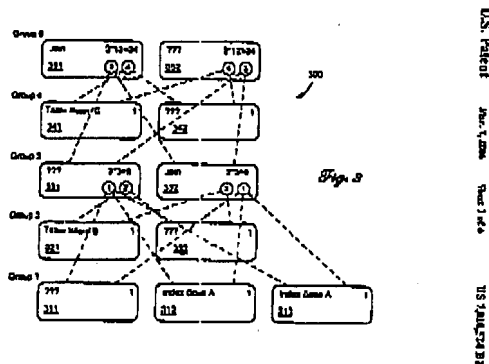
In addition, with regard to the limitations "generating a query execution plan with a super group block having an array of pointers, wherein each pointer points to a linked list representing grouping sets for a particular level," these limitations are not described by Galindo-Legaria at col. 5, lines 25-34, col. 5, lines 56-63 and Figure 3, which are set forth below:

Galindo-Legaria: Col. 5, Lines 25-34

The optimizer stores the alternative plans in data structure 250. This data structure, a table in this embodiment, stores alternative operations and their interconnections at a number of different levels, as described hereinafter. This table is not destroyed in the process of determining an optimum plan, but is kept for later construction of alternative plans other than the single optimum plan. Ranking module 260 builds a directory 251 having pointers to the locations of various operators and groups within structure 250, and computes "rank data."

Galindo-Legaria: Col. 5, Lines 56-63

FIG. 3 is a symbolic diagram 300 of a portion of an illustrative table 250 for compactly encoding information required to construct multiple execution plans so as to take advantage of the many common parts among different alternative plans. A number of groups, five in this example, each contain a collection of operators that point to other groups as children. Each candidate plan is a tree of these operators extracted from the groups.

Galindo-Legaria: FIG. 3

The description from Galindo-Legaria set forth above describes a structure to store alternative query execution plans, but the only “groups” described are groups of operators that are shared among alternative query execution plans.

However, this structure in Galindo-Legaria has nothing to do with “grouping sets comprised of grouping columns” as recited in Applicants’ claims.

Instead, the structure of Galindo-Legaria has a completely different purpose, namely grouping common operators shared among different alternative plans.

Nonetheless, the Office Action suggests replacing the groups of operators in Galindo-Legaria with the grouping sets comprised of grouping columns in Cochrane.

However, such a modification would change the operation of Galindo-Legaria; indeed, it would render Galindo-Legaria inoperative, because the structure of Galindo-Legaria would comprise grouping sets (which are not operators) rather than operators.

Thus, Applicants’ attorney submits that independent claims 1, 4 and 7 are allowable over the combination of Cochrane and Galindo-Legaria. Further, dependent claims 2, 3, 5, 6, 8 and 9 are submitted to be allowable over the combination of Cochrane and Galindo-Legaria in the same manner, because they are dependent on independent claims 1, 4, and 7, respectively, and thus

OCT 25 2007

contain all the limitations of the independent claims. In addition, dependent claims 2, 3, 5, 6, 8 and 9 recite additional novel elements not shown by the combination of Cochrane and Galindo-Legaria.

III. Conclusion

In view of the above, it is submitted that this application is now in good order for allowance and such allowance is respectfully solicited. Should the Examiner believe minor matters still remain that can be resolved in a telephone interview, the Examiner is urged to call Applicants' undersigned attorney.

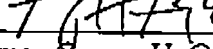
Respectfully submitted,

GATES & COOPER LLP
Attorneys for Applicants

Howard Hughes Center
6701 Center Drive West, Suite 1050
Los Angeles, California 90045
(310) 641-8797

Date: October 25, 2007

GHG/

By: 
Name: George H. Gates
Reg. No.: 33,500

G&C 30571.291-US-01